# The National Computational Science Alliance Access Grid: An Internet-Based Collaboration Tool Augmented by High-Performance Computing

Robert A. Ballance *
Thomas P. Caudell †
John Greenfield‡
Rick Stevens§

June 2, 2000

## Abstract

The National Computational Science Alliance Access Grid is an Internet-based conferencing system that supports real-time, multi-point, group-to-group communication and collaboration. The Access Grid provides a collaborative environment that complements the Alliance's Computational Grid, for remote visualization, application development, planning sessions, and other events or meetings. The Access Grid is also being adapted for distance education and telemedicine. Based on concepts developed at Argonne National Laboratory, the Access Grid is the result of collaborations among Alliance member institutions including Argonne National Laboratory, the University of New Mexico, the University of Kentucky, Boston University, and the University of Utah.

## 1 Introduction

All good systems reflect one or more fundamental metaphors in their design. For the Macintosh, it is the "desktop;" for Microsoft Excel (and all spreadsheets) it is the table of data; for the National Computational Science Alliance (Alliance) Access Grid (AG), it is a café where the regulars meet to exchange ideas[12]. This is a familiar setting, akin to the "water cooler." The difference is that with the Access Grid, the participants may be separated in space but not time. Some of the "locals" may be physically located in Illinois, others in New Mexico, still others in Kentucky or Massachusetts or Hawaii. The "Café Model" stresses informal interaction in a *persistent, designed virtual space.* The Access Grid, of course, can also support formal interaction techniques as well, such as the National Chautaquas [3] in 1999 and 2000. The Chautaquas are large-scale, multi-site, multi-participant meetings held by the entire Alliance Access Grid community.

Technologically, the Access Grid is an Internet-based conferencing system that supports real-time, multi-point, group-to-group communication and collaboration. The Access Grid provides a collaborative environment that complements the Alliance's Computational Grid[7]. The Computational Grid links together machine resources: computers, storage, and software for use in scientific programming. The Access grid links scientists, application users, and developers by providing remote visualization, ap-

*Albuquerque High-Performance Computing Center (AHPCC) & UNM Dept. of Computer Science.
†UNM Dept. of Electrical Eng. and Computer Engineering & AHPCC.
‡AHPCC.
§Argonne National Laboratories & Dept. of Computer Science, University of Chicago

plication development, planning sessions, and other events or meetings. The Access Grid is also being adapted for distance education and telemedicine.

Access Grid studios are operational from the East Coast to Hawaii, including studios at the Albuquerque High-Performance Computing Center (AHPCC), and the Maui High-Performance Computing Center (MHPCC), a DoD Shared Resource Center. Table 1 lists the current sites. Studios at Air Force Research Laboratory sites (Kirtland AFB and Hanscom AFB) are slated for deployment during 3Q00 during a preliminary evaluation program for Access Grid technology. If the technology proves viable within AFRL, other studios throughout AFRL and DoD may be deployed. Temporary studios have successfully been set up in Europe (Moscow, Paris), on the display floor at SuperComputing '99, and in Tsaile and Shiprock New Mexico. The latter studios are located within Dineh College, a tribal college operated by the Navaho Nation. In this paper, we present a snapshot of the Access Grid as of Spring 2000.

Befitting a distributed, collaborative environment, the Access Grid itself is the product of a distributed, collaborative team. The insight and the vision for the AG orgininated with Rick Stevens' Futures Lab at Argonne National Laboratory (ANL). ANL continues to lead in the development and deployment of revised as well as newly-introduced Grid features. Other AG contributors include the Boston University, the University of Kentucky, the University of Utah, and the University of New Mexico through both the AHPCC and the Maui High-Performance Computing Center (MHPCC). More recently, an Access Grid deployment effort has begun the rollout of Access Grid connection points (studios) throughout the Alliance. The deployment effort is being lead by the Ohio Supercomputing Center. The authors of this paper, while participants in the AG effort, view themselves also as reporters speaking on behalf of the entire Access Grid Team.

# 2 The Access Grid

Access Grid users share audio, video, presentations, visualization environments, browsers, whiteboards and graphics tablets, and even remote sensors and devices. Early experience indicates that the Access Grid provides a high-bandwidth collaboration venue in which participants can freely and informally interact.

Grid studios, the access points, typically combine large-screen multimedia displays using conventional projectors with economical high-quality audio. Studios scale from individual offices to auditoriums, with the most effective size being a small to mid-size conference room. By using Open Source software and COTS components, the cost of individual studios is kept low. The Access Grid substrate is the high-bandwidth Internet, using IP multicast and middleware to feed the sites.

The Access Grid itself consists of a number of physical *studios* physically linked via the high-speed Internet (vBNS, Abilene) and logically linked by Open Source and commercial software using IP multicast protocols. Several external forces shaped our selection of technologies:

1. The Access Grid project originated as the development of a research platform into collaborative computing. As such, the original participants attempted to reduce costs by employing COTs hardware and software.

2. The Access Grid developers are committed to using Open Source software [13] wherever it is available and appropriate. This goal is consistent with the well-chronicled development of Internet protocols and systems. Open Source also allows the various development teams to freely share their code; a freedom essential to collaboration.

3. The cost of high-bandwidth communication is falling rapidly, while the cost of commercial video conferencing equipment is still relatively high. Our goal is to design a system with entry costs for a single studio in the range of $\$40 - \$50,000$.

| | |
|---|---|
| ACCESS DC | Argonne National Laboratory |
| Boston University | Lawrence Berkeley National Laboratory |
| Kansas University | Los Alamos National Laboratory |
| The National Center for Supercomputing Applications | |
| North Dakota State University | Ohio Supercomputing Center |
| Princeton University | San Diego Supercomputing Center |
| University of Illinois at Chicago, Electronics Visualiation Lab | |
| University of Kentucky | University of Utah |
| University of New Mexico, AHPCC | University of New Mexico, MHPCC |

Table 1: Access Grid Studios, May 2000

4. Native IP multicast [14] is becoming widely available. With multicast, a number of sites can share data streams within their own group, without requiring centralized servers or hardware support beyond the router. Multicast, therefore, is a natural underpinning for IP-based conferencing.

5. A larger fraction of an individual researcher's time is spent in conversation with coworkers and collaborators spread across the globe. This is a function both of the decreased cost of collaboration (due to the Internet and decreasing costs of bandwidth) and of the increased size and scope of projects. For example, several scientific application communities are focussing their efforts into shared "community" codes such as Cactus [6] .

6. Participation in an Access Grid meeting should be as low-overhead as possible. This means both operational simplicity, and user simplicity. A participant should be able to join a meeting just by walking into the meeting space.

These forces led the Futures Laboratory at ANL to the following vision: *The Access Grid should suport multi-person, multi-site collaboration in an informal, persistent, immersive environment that is cost-effective for individual sites, that is based on commercial hardware, and that uses easily obtainable, modifiable, and sharable software.*

## 2.1 Access Grid Studios

Access Grid studios are built using commercial off-the-shelf (COTS) hardware coupled to Open Source or easily obtainable software. The problem of equipping a studio is therefore largely moot. Any organization with the financial and technical wherewithal can deploy a studio on its own. Today, most organizations currently rely heavily on existing Grid users to assist them in their deployment effort. However, the Alliance/AG deployment team is working on ways to simplify the construction, tuning, and operation of the studios.

The studio itself constitutes a context in which your presence indicates your agreement to participate. Thus an individual controls their own privacy by controlling their physical location.

Access Grid studios come in three basic sizes:

**Café** The café studio is the primary designpoint for Access Grid nodes. It is designed to support 5–15 participants per studio in a physical space of roughly 200–400 square feet. This is the size of an average conference room.

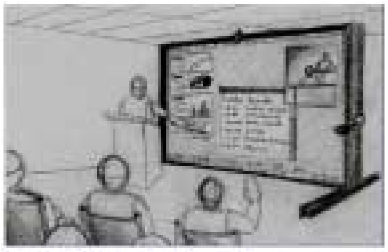

Figure 1: Studio Layouts

**Office** The office studio is a down-scale version of the Café, and is more akin to a conventional video conferencing system (single computer, single monitor, single camera). Experiments with office-scale studios are just beginning.

**Theater** Scaling up, the studio can be embedded into a larger space. During the first Chautauqua for the Access Grid, held at UNM last August, the studio was embedded into a large auditorium that held over 100 attendees. This experiment showed that the AG studio can be scaled up for large meetings. However, given the scale of the meeting, the actual content of the meeting had to be more formal (with trained emcees) in order to keep the proceedings on track.

In the remainder of this report, we shall focus on the café-sized studio.

## 2.2 The Café Studio

A typical café studio consists of a designed room which transmits several (2–4) video streams plus a public audio stream to the other studios in the Grid. Within the studio, the incoming video streams are displayed on one or more large (wall-sized) screens. This is accomplished by using multiple LCD display projectors, each able to display 1, 4, or 9 separate video streams. In the café sized studio, 2 or 3 projectors are configured to provide up to 27 separate windows. Figure 1 illustrates two possible studio layouts used by Access Grid sites.

Each studio controls both its outputs to the other studios, and the layout and display of incoming audio and video streams. Currently, the layout and selection of streams is left to a human operator at each studio. However, work is underway to develop the software necessary to allow remote control of the studios. On a related tack, we also envision developing software which can allow local participants to see what the operator sees—that is, the separate streams coming into the studio. This "personal studio" will enable individual participants to control their own local views. This, of course, is not without its problems in a collaborative environment; witness both the usefulness and the distraction of laptop computers in any large meeting today!

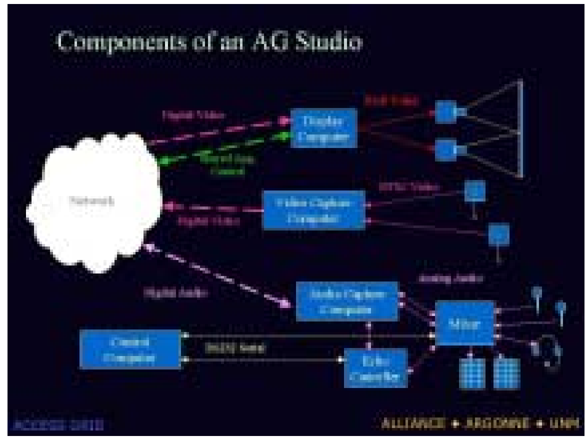


Figure 2: Studio Architecture

### 2.2.1 Café Studio Hardware

The café-sized studio uses four major audio and video components to provide an immersive conferencing environment. The entire system is driven by four standard PC computers. Currently, the computers in a studio run either Windows NT or Linux, depending on their roles. Figure 2 provides a schematic layout of the hardware configuration.

**Video capture and transmission** Multiple video feeds are extremely desirable due to their ability to show more than one view of any given room. The camera configuration uses

- 1 camera directed at the speaker (the *presenter camera*)

- 2 or more cameras focussed on either the audience or the local video display.
- Optionally, a document camera for transmission of document images.

In use, the audience cameras allow other participants to look around the meeting room, while the local video display camera gives other rooms to the chance to see what a given studio is seeing. The latter is necessary because each studio currently controls its own screen contents. The video camera inputs are connected to a video capture computer which transmits the video streams into the network.

**Video reception and display** Multiple large screen (mural) displays are highly desirable, due to the number of video streams that can be displayed and to the larger size of the displayed images. Projected displays also seem to enhance the user's perception of presence within a meeting. Our conjecture is that this is due to the user's vision being focussed outwards, toward participants rather than being focussed more narrowly inwards toward a conventional monitor.

The displays are provided by LCD projectors driven by multiple display video boards on a Windows NT or Windows 2000 machine. This allows all of the displays to function as a single very large desktop, with an additional monitor used as the left-most side of the desktop. The monitor side of the desktop contains the toolbar and desktop icons. Windows are placed onto the projected display simply by dragging them from the desktop onto the screen. When the projector's images are carefully aligned side by side, a two or three projector wide empty seamless single desktop can be achieved. On the projector portion of the desktop, the operator can place video streams in large medium or small sizes as well as any PowerPoint, Web browsers, or other window displays that are desired by the viewers.

In use the large display allows the participants to see multiple views of other participants from several different locations in near life size.

**Audio capture and transmission** Studios are currently configured using a variety of microphone options. Ambient mikes seem to be the most desirable, but are the hardest to configure and adjust. Wireless hand-held and lavalier microphones are effective, but counter our design goal of minimal overhead to participate in a session. Table-mounted microphones have been used, but have the drawback of limited pick-up range. However, table microphones are effective at picking up the key clicks from a nearby laptop user. During studio construction, headsets have also been used. Whatever miking system is chosen, the microphone inputs are mixed into an outgoing audio stream.

Echo cancellation is required to prevent a site from retransmitting everything it "hears" back to all the participants. While a site without echo cancellation will not notice any problems every other site will hear everything they and all other sites say echoed after a fairly brief delay. This echo quickly renders any conversation impossible.

**Audio reception and playback** Audio playback is handled by 2 or more good, quality speakers. Both audio reception and transmission are handled by an off the shelf sound card on a single computer, separate from the computers handling video and display.

A single PC can handle both audio capture and playback. However, currently a second Windows PC is used to boot and control the echo cancellation gear.

### 2.2.2 Café Studio Software

The Access Grid software uses recent editions of the MBONE conferencing applictions [11]. The VIdeo Conferencing application (VIC) is used to transmit and receive video stream. The Rational Audio Tool (RAT) is used for audio transmission and playback. The Access Grid also uses a number of locally- developed tools, such as the Virtual Venue software for handling different conferences within the AG. Almost all of the Grid's software runs on either Microsoft Windows or Linux platforms.

VIC is a multicast-enabled video conferencing tool developed by the Network Research Group at Lawrence Berkeley National Laboratory, in collaboration with the University of California, Berkeley. It supports the H.261, H.263, and H.263+ video codecs and software JPEG compression.

RAT is a multicast-enabled Open Source audio tool. Like VIC, RAT is part of the MBONE Conferencing suite. The most recent editions or VIC and RAT support synchronization between the two applications via the "conference bus."

Distributed Powerpoint is a controller for Microsoft PowerPoint$^{\mathrm{TM}}$ so that a single user can control the display of PowerPoint slides at multiple locations.

#### 2.2.3 Studio Space Requirements

In our vision, the Access Grid relies on carefully designed spaces as nodes. In the best scenario, projectors are permanently mounted with the correct merging of images, microphones easily pick up the voices of particpants (but not the background noise of the air conditioning), and the computers are unobtrusive. Over the past year, we have come to realize that we are building a very special form of broadcast studio: one that combines the lighting, sound, and camera quality of broadcast with the ambiance of a sitting area. The tools of any good meeting: wall charts, whiteboards, A/V componentry, and laptop connections need to be close at hand and easily managed.

Overall, each site's studio differs from all the others. At ANL, there is a studio using projection onto two walls, with the speaker in the corner. One wall is used for the audience, the other for the presentation. At UNM and ANL, we are using whiteboards that can double as projection screens so that the two media can easily be combined.

Comfort and ease of use for the meeting participants cannot be over-stressed. For a collaboration to succeed and mature, the participants cannot be handicapped by the media. Participants in AG sessions agree that with sufficient bandwidth, the effect of the Access Grid is to encourage useful collaboration.

### 2.3 Access Grid Networking

Obtaining and installing studio equipment is the "easy" part. Harder is the problem of obtaining sufficient effective network access. The Access Grid is a cheerful consumer of the most bandwidth it can preempt. The initial Chautauqua meetings provided an opportunity to truly test (and subsequently correct) the performance of the high-bandwidth Internet.

#### 2.3.1 Bandwidth

A single Access Grid studio typically transmits 2-4 video streams, an audio stream, and (potentially) several other media streams with relatively lower bandwidth requirements. Each video stream requires an effective bandwidth of 128-512Kbs (or more), and an audio stream requires comparable bandwidth. The total bandwidth required throughout the network for a single meeting session is therefore the number of active studios multiplied by the bandwidth transmitted by a single studio. 10Mbs of effective bandwidth is not enough.

Inadequate networking support results in lost or delayed packets. This shows up as choppy video or unintelligible speech. Preliminary experience indicates that poor video quality is tolerated by participants far more than audio problems. The bandwidth required is true, effective bandwidth, not the nominal rating of the home cable. For example, a 10Mbs ethernet LAN typically delivers a sustained effective bandwidth of only 3-4Mbs. Second, the bandwidth must be delivered end-to-end. If you have a Gigabit ethernet connection to your nearest backbone, but the studio that you are talking to has only a 10Mbs connection, you get only the effective bandwidth of the 10Mbs connection. Finally, consider the the intervening routing and links. A temporary studio at the Southwest Indian Poly technical Institute (SIPI) in Albuquerque was unable to effectively communicate with a studio at AHPCC (15 minutes apart by road) due to routing that sent packets from New Mexico to Washington, DC and back—via California!

### 2.3.2 Multicast

Native IP multicast [14] uses UDP datagrams to "broadcast" data packets to a select group of receivers. Multicast routers employ sophisticated techniques to minimize the duplication of packets until absolutely necessary. For instance, if a computer in Boston is broadcasting multicast data to three sites on the West Coast, all three of which are served by the same router in Los Angeles, only a single packet will travel transcontinentally. At the LA router, the packet will then be split out into three copies, one for each local machine. This technique minimizes network bandwidth usage in the backbones. However, the use of UDP datagrams means that the protocol is unreliable; packets *may* be lost without opportunity for recovery. Fortunately, with voice and video data, occasional packet loss can be tolerated.

The presence of multicast groups, all listening to the same "channel" provides a useful structure for isolating conversations and for developing privacy mechanisms. The Virtual Venue software provided by ANL to the Access Grid uses the multicast groups to provide separate conferencing channels within the Grid. Each venue acts like a channel selector. Virtual venues therefore allow multiple simultaneous conversations to take place within the Access Grid.

Access Grid software requires either the use of multicast-capable routers or the use of a traffic reflector such as the Fermi MultiSession Bridge currently in place at ANL. A traffic reflector provides a means around lack of multicast, at the cost of increased traffic to and from the bridge.

Multicast is a relatively new addition to the Internet; many providers of high-speed bandwidth provide limited support for its features. Vendors with working systems tend to be slow to upgrade to provide new features, especially if the requestor is down stream and not an immediate customer. There may be a transition period in which Access Grid vendors and support teams have to work as closely with network providers as with the organizations installing studios.

Another stumbling block with multicast is that not all vendors implement multicast the same way. Cisco appears to be the leader in developing the multicast standard, but not all manufacturers track the standard closely, and not all network providers want to commit to Cisco. The problem is reminiscent of Microsoft vs. Linux in the operating system world.

## 3 Access Grid Research

Research and development on Access Grid applications and improvements is underway at many of the partner sites. Adding scientific visualization into the mix of data streams is a high priority for the AG team since an integrated visualization capability will tie the Access Grid closely to the Computational Grid. Section 3.3 discusses one approach to integrating visualization with the Access Grid.

We also imagine entering an AG space with personal devices, from cell phone and pager to laptop computers, and having the devices automatically integrated into the AG. We envision a method of exporting or importing content from these devices, either to or from other individuals, or groups of individuals to the AG session participants as a whole.

For structured events, a script is required. For an AG event, that script should be a tool to automate and synchronize each activity for each event across the network. The tool should be able to verify that required resources are in place at each site and marshal them if not. The tool should launch applications, start media players, remind player, request and schedule resources and generally automate the performance of a network based structured meeting.

The above list does not exhaust the research opportunities inherent in the AG. Other areas, which are currently being investigated but which are not sufficiently developed to elaborate further are: (1) Human Interface issues; (2) Non-traditional input and output device integration; (3) High resolution Visualization; (4) Streaming Media recording and playback; (5) Integration with other standards (H.323);(6) Collaborative tools; and (7) White Boards, web browsing, document server, application sharing.

This section briefly reviews the projects and status of several research projects underway at UNM.

## 3.1 Audio and Video Quality

Improvement of both the audio and video quality of the Access Grid is necessary. The AG does not yet meet the standards set by more expensive, commercial videoconferencing systems. Many of these systems are based on the H.323 [8]. Several of the Access Grid sites are working on a variety of mechanisms to improve the overall quality of the Access Grid.

There are two primary issues with the current audio and video configuration. First, the audio and video streams are sent separately through the Internet. This means that packets containing voice data may arrive before or after the packets containing the corresponding video data. When the two streams are slightly out of sychronization, this causes apparent inconsistencies between lips and sound. With more network delays, this separation can lead to lags and dropouts. Second, the current studio architecture places the audio and video processing on different machines, making synchronization more difficult.

There are three possible improvements. The first is to better synchronize the streams. RAT and Vic already provide synchronization facilities when they run on the same machine. Development is currently underway to implement cross-machine proxies that will enable synchronization between processors.

Second, the current system can be enhanced to use more complex adaptive algorithms that will minimize the chance of packet loss and maximize the delivery of data across an unreliable network. This research leads into more sophisticated network transmission algorithms and new or revised compressor/decompressor (codec) algorithms.

Third, we are beginning to look at algorithms that combine the audio and video into a single stream. One such algorithm is MPEG-1. MPEG-1 hardware encoders are still relatively expensive, placing them outside the current price point for Access Grid studios. However, as prices drop, it may be possible to replace the current software codecs with better hardware codecs.

Finally, as the Access Grid expands, it is being deployed into environments that already use commercial videoconferencing equipment. Also, current users of such equipment within the Alliance are talking about ways to connect H.323 systems to the multicast environment. While such users would not have the same immersive environment that was the primary goal of the AG, they might have limited connectivity into AG conferences.

## 3.2 3D sound localization

With as many as 27 video images on the screen it can be difficult to determine which video ontains the person currently speaking. To rectify this problem, we are experimenting with the use of 3D localized sound to make it appear that the sound is coming from the image which contains the speaker. We are currently evaluting the effectiveness and usefulness of currently available COTS systems for making sound appear to come from displayed images. If that research proves successful, we expect to start the work of integrating the sound localization into the Access Grid.

A second development project at UNM is addressing the same problem by trying to determine which window on the screen best shows the current speaker, and then highlighting the window. In the first version of the program, the operator on the transmitting end of the sound stream will select the appropriate window, and the video and audio streams will then coordinate their display. This coordination is also necessary for proper sound localization.

With localized sound, there is a problem of where to place the audio from a speaker not currently visible. One approach is to use a speaker placed in the rear of the room. This would allow us to simulate the "questions from the back of the room" that occur in localized meetings.

## 3.3 Flatland on the Grid

Visualization of data and simulations is a crucial part of the scientific process. The AG offers teams an effective work venue, but does not directly provide them with access to their data. This issue is addressed in a research project using the UNM visualization/virtual reality tool called Flatland [1]. Two technical issues must be resolved for this tool to be used effectively: 1) graphical information from a visualization machine must be extracted from the visualization software and transferred to a AG video stream, and 2) user interactions must be transferred from the AG environment into the visualization software. The Flatland system is serving as a testbed for studies in both of these areas.

Flatland is portable cross computing platforms such as Unix/Linux/NT. It uses a standard graphics language (OpenGL or MESA) and associated libraries for its implementation of graphics. Flatland has been designed to support a wide sensorium, including 3D graphics, 3D sound, simple haptics, and verbal conversations, both with artificial elements in the environment and other immersed workers. Each user operates from within their personal workspace represented as a surrounding dodecahedral vessel, which is flown through Flatland and contains customizable consoles, screens, and controls, as well as their personal tool chest. The following section gives a more detailed description of this system.

### 3.3.1 Description of Flatland

Flatland allows software authors to construct and users to interact with arbitrarily complex graphical and aural representations of anything that can be programmed. Flatland is written in C++ and uses the standard OpenGL graphics language extensions to produce the 3D graphics. In addition, Flatland uses the standard GLUT library for window, mouse, and keyboard management. Flatland is multithreaded and uses dynamically linked libraries (DLL) to load applications that construct or modify its virtual environment (VE).

At the core of Flatland is an open, custom Translational Graph data structure that maintains and potentially animates the geometric relationships between the Objects in the Graph. Graph Objects contain all of the information necessary to draw, sound, and control the entity represented by the Object. The Graph is one part of an higher level structure referred to in Flatland as a Universe. The Universe contains the scene Graph, a flat database of Objects in the Graph, and a reference to the Graph vertex that is currently acting as the root of a hierarchically organized tree structure. The tree is the Graph with one node lifted into the role of a tree root node. Normally this root node is used as the camera for rendering. All other Graph nodes hang below, forming a tree structure. Flatland maintains the geometric transforms as the Graph is reconfigured into different tree instantiations. The tree is a hierarchy in the sense that the children inherit and utilize the geometric transforms of the parent. There are no other semantics related to the position of an Object in a tree instantiation. For example, there are no intrinsic levels of abstraction related to the position of an Object in a tree.

Flatland is intrinsically multi-threaded, allowing the system to make use of computer systems with multiprocessors and shared memory. The main thread spawns an OpenGL graphics thread, a Flatland sound thread, and a real-time tracker thread. The optional tracker allows the user to use the 3D interaction metaphors with their applications and to use head tracking or eyepoint motion tracking.

An application in the context of Flatland is a relatively self contained collection of objects, functions and data that can be dynamically loaded (and unloaded) into the Graph of an environment. An application is responsible for creating and attaching its objects to the Graph, and for supplying all object functionality. An application is added to Flatland through the use of a Configuration file. This structured file is read and parsed when Flatland starts, and contains the name and location of the DLL that has been created for the application, as well as a formal list of parameters and an arbitrary set of arguments for the application. The Flatland configuration file distinguishes two types of applications: 1) base apps and 2) user apps. Base apps are loaded at the startup of Flatland, while user apps are not. User apps may be loaded by several methods, including the use of buttons in the VE or standard pull down menus.

In Flatland, graphics and sound are treated symmetrically. Each object in the Graph contains, among other things, a draw function and a sound function. The draw function contains or calls all of the code to draw and animate the graphics that represents the object. From an author's perspective, all object graphics is based on and drawn its own model coordinate system. Other structures in the Graph handle the placement and orientation of the object's model coordinate relative to other objects in the Graph and subsequently the camera. The graphics thread loops over all of the objects in the system, sets the appropriate transformation matrices, and calls their draw functions.

The sound function within an object contains all of the calls or code to make sounds that represents the object. Flatland has a library of sound function calls that are similar to OpenGL. Wave sound files are treated like OpenGL display lists and are called sound lists. In addition to opens a sound list, functions exist that allow the author to control the starting, looping, stopping, volume, and 3D location of the sound. All sound is emitted in Flatland from point sources. The author specifies the location of the sounds in the same model coordinate system used for graphics. The sound thread loops over all of the objects in the system, sets the appropriate transformation matrices, and calls their sound functions. Currently no online sound synthesis is available in the sound libraries, although this capability is under development.

Although real-time six degree of freedom trackers are not generally available, Flatland can make use of them when they are. A tracker is a multiple degree of freedom measurement device that can in real-time monitor the position and/or orientation of multiple receiver devices in space relative to a transmitter device of some sort. As such, Flatland launches a tracker thread to sample the available tracker information and make it available for use by applications. In the standard Flatland configuration, trackers are used to locate hand held wands and to track the position of the user's head. Head position and orientation is needed in cases that involve the use of head mounted displays or stereo shutter glasses. Commercial trackers typically supply data at rates of 30–60 Hertz.

User interaction is a central component of Flatland. Two general classes of interaction are offered: from within the VE and from without. By "within" is meant from within the 3D environment where the user uses devices to interact with objects in the 3D environment. Flatland has several types of "within" interaction modalities., including the wands to select objects and to activate buttons located on objects, voice recognition system that permits the user to speak commands to an object, and speech synthesis that allows the objects to talk back. From "without" is meant standard menus, buttons and sliders in control panels drawn in separate windows. These standard interactors are provided through the GLUT libraries and a custom widget library. Each object in the environment could potentially have an external control panel.

Source code, example applications and documentation are available on the our web page at www.ahpcc.unm.edu/homunculus.

### 3.4 Transport of Compressed Images and Videos

As mentioned above, graphically generated video must be piped into a AG video stream for scientific visualization tools to be effective. This will involve video compression methods to increase the graphical frame rate. Different compression methods can be used for images and videos [10]. The current Access Grid uses the standard H.261 package which is good at video conferencing and multicasting. MPEG gains popularity as well, especially for its richness in its scalable base stream and other augment substreams, in its well defined transport formats, and in its object-oriented approach. We are conducting experiments for the JPEG and MPEG standard as an alternative for the comparison purpose.

MPEG-4 is designed for a multimedia system with interoperable communication of complex scenes including audio, video, synthetic audio, and graphic material. The major advantage of MPEG-4 is its extremely low bandwidth requirement, making it suitable for Internet and wireless transmission. Particularly, MPEG-4 is able to compress graphics materials very well. We are developing an MPEG-4 system to test the trasfer of graphical images onto the Access Grid.

For the current stage, integrating a graphical image stream as a special stream is an immediate addition to the Access Grid. Without developing rendering capacity at the user site, we may use our existing MPEG-2 system to compress images to MPEG-2 and transmit it on the Access grid. We are investigating more details, such as tradeoff of image loss and degree of compression, availability of limited interactive control, and frame rate of transferred video streams. A graphics to MPEG-4 translator is being developed to convert a graphics format to very low bit rate stream MPEG-4 [9].

### 3.5 TOUCH Telemedicine Project

The Telehealth Outreach for Unified Community Health (TOUCH) project [4] is using AG technology to demonstrate the feasibility of employing advanced computing methods to enhance education in a problem-based learning (PBL) format currently being used in a medical school curriculum, applying a specific clinical case (traumatic brain-injury) as a model, and deploying to remote sites/workstations. The John A. Burns School of Medicine and The University of New Mexicos School of Medicine have chosen to collaborate on this special project. Hawaii and New Mexico face similar healthcare challenges in providing and delivering services and training to remote and rural areas. Both states must deal with common challenges such as barriers to healthcare access (water in Hawaii, land in New Mexico) unique indigenous populations, large multicultural populations, and isolation of healthcare professionals and students/trainees in remote settings. Recognizing that health care needs are local and require local solutions, both states are focused on improving health care delivery to their unique populations and have begun to benefit from sharing information and experiences. Emerging telehealth technologies can be applied using existing high performance computing and communications resources present in both states.

Phase I of this project is aimed to improve the quality of health care service and education in remote, multicultural areas in Hawaii and New Mexico. This will be accomplished through modern telehealth technologies using advanced information systems and high performance computing in order to enhance education, training, patient care management and problem solving in collaboration with students and health care providers at dispersed locations. The schools of medicine in both states, in collaboration with their rural hospitals and training sites, propose using telehealth technologies to enhance and deploy their existing problem-based learning (PBL) curricula in the training of students and health care providers from a variety of disciplines to medical practice in settings serving the unique health care needs of the diverse populations.

In order to achieve these objectives, these schools and two of their respective training sites (Northern Navajo Medical Center and Maui Community College) with support from their local high performance computing centers, the Albuquerque High Performance Computing Center (UNM HPCERC) and the Maui High Performance Computing Center (MHPCC), both centers of The University of New Mexicos High Performance Computing Education and Research Center (UNM HPCERC), will integrate the use of telehealth technologies, including the AG and Flatland, into the curriculum on brain injury and into ongoing education opportunities. This collaborative program development between the two states, teaming healthcare professionals, educators, librarians and students with computational scientists and engineers, will serve as a model for training in other rural environments.

### 3.6 Security and Privacy

Right now, the Access Grid is basically a party line; any studio can connect at any time. This works well to meet the original goal of supporting a "cafe" style of interaction. The Virtual Venue system deployed during SC99 allows the users of the Grid to divide the available network into separate channels (venues) so that their conversations don't overlap, but there is currently no ability to either (i) completely determine which studio is on the line or (ii) keep conversations private.

This is problematic for academic users who are pursuing possible intellectual property protection for their research, as well as for military and business users. The basic audio and video processing software (RAT and VIC) are designed to support low-level encryption of their data streams, but there is still a need to actually select, implement and deploy that functionality.

Once the data streams can be encrypted, there is still a need for authentication and key management. The Public Key Infrastructure [2] (PKI) facilities currently available through Globus [5] may provide the necessary functionality at low overhead. This is certainly a valid approach for the Grid community since the use of Globus can be either assumed or required. There may also be a need for directory support; this is also a feature in Globus.

A project to combine the Globus PKI infrastructure with the RAT/VIC streams to support both authentication and encryption is being initiated this summer at UNM.

## 4 Results

While meeting its original design goals, the Access Grid is poised to become a victim of its own success. An attractive solution to teleconferencing and collaboration, the Grid is also stressing the limits of high-bandwidth network communication. The Grid community needs to continue work closely with both funding agencies and network suppliers to assure that the high-speed, multicast network is rapidly deployed.

There is much work to be done in making the Access Grid truly accessible to ordinary consumers of high-technology. This work include both the assurance of stability as well as the addition of features that will contribute to ease of use.

By founding the Access Grid on easily obtained hardware and Open Source software, the number of researchers and developers who can contribute to the project is largely unlimited. Access Grid steering groups, including the Alliance Deployment Team and the recently-formed AG Engineering Task Force are taking up the role of coordinating projects and steering the overall roll-out the AG. You can help by joining a project, contributing to a steering team, or even developing new applications.

## 5 Acknowledgements

Along with the entire Access Grid team, we wish to acknowldege Dr. Min-You Wu, Dept. of EECE, UNM and Dr. Dale Alverson, Dept of Pediatrics, UNM Health Sciences Center.

## References

[1] A description of flatland, open source, and a copy of the user's manual. www.ahpcc.unm.edu/homunculus, 2000.

[2] Carlisle Adams and Steve Lloyd. *Understanding Public-Key Infrastructure: Concepts, Standards, and Deployment Considerations.* Macmillan Technical Publishing, 1999.

[3] Alliance chautauqua 200. http://www.ncsa.uiuc.edu/alliance/-chautauqua/, 2000.

[4] D. Alverson. Telehealth outreach for unified community health (TOUCH). National Institutes of Health Office of Telemedicine, August 2000.

[5] Randy Butler, Doug Engert, Ian Foster, Steven Tuecke, John Volmer, and Von Welch. Design and deployment of a national-scale authentication infrastructure. Technical report.

[6] The cactus code server. http://www.cactuscode.org/, 2000.

[7] Ian Foster and Carl Kesselman, editors. *The GRID: Blueprint for a New Computing Infrastructure.* Morgan Kaufmann Publishers, Inc, 1999.

[8] Olivier Hersent, David Gurle, and Jean-Pierre Petit. *IP Telephony: Packet-based Multimedia Communications Systems.* Addison-Wesley, 2000.

[9] C.S. Lin, W. Shu, and M.Y. Wu. Transmitting variable-bit-rate videos on clustered vod system. In *IEEE International Conference on Multimedia and Expo (ICME2000)*, July 2000.

[10] C.S. Lin, X. Zeng, M.Y. Wu, and W. Shu. Vod on internet 2. In *International Conference on Computers and Their Applications (CATA-2000)*, March 2000.

[11] Mbone conferencing applications. http://www-mice.cs.ucl.ac.uk/multimedia/-software/.

[12] Ray Oldenburg. *The Great Good Place : Cafes, Coffee Shops, Bookstores, Bars, Hair Salons, and Other Hangouts at the Heart of a Community.* Marlowe & Co, 3rd edition, August 1999.

[13] Eric S. Raymond. *The Cathedral and the Bazaar.* O'Reilly & Associates, 1999.

[14] Beau Williamson. *Developing IP Multicast Networks.* Cisco Press, October 1999.